Increasing the Availability of the Computerized Patient Record

Dario A. Giuse, Andrew Mickish.

Informatics Center, Vanderbilt University Medical Center, Nashville, TN.

The MARS clinical repository, originally developed at the University of Pittsburgh, provides electronic access to the patient record at Vanderbilt University Medical Center. The original client interface we developed runs on all standard clinical workstations in the medical center, but is operating-system dependent. Porting and maintaining it on the variety of hardware- and software combinations found on VUMC personal computers would be fairly costly. To broaden the availability of the system to faculty and health care providers in all areas, and to support future access from Vanderbilt-affiliated providers outside the main campus, we are developing a new Web-based client. The new client provides good functionality and performance, and will be a strategic asset in our long-term commitment to making relevant clinical information immediately available to authorized health care providers.

INTRODUCTION

The computerized patient record (CPR) has been a key addition to health care practice in the past three decades.[1-5] In addition to increasing the speed and reliability with which clinical information about individual patients can be made available to care providers, the CPR also a) supports clinical research, by allowing clinical data about patients with given characteristics to be collected easily;[6,7] b) makes it possible to generate reminders for care providers;[8,9] and c) may improve the utilization of resources such as lab tests and drug prescriptions.[10,11] Access to the CPR has traditionally been within a single geographical area, such as an individual hospital or medical center. However, the CPR could become a fundamental infrastructure in the changing world of health care, by facilitating information sharing across different organizations that occasionally provide care for the same patient. Incompatible information systems at different organizations have usually made this idea difficult to implement.

This paper describes the Vanderbilt University Medical Center (VUMC) strategy for expanding access to the MARS electronic patient record system, which stores all clinical inpatient and outpatient information, beyond the current single-platform environment. The strategy uses the World Wide Web as the interface, providing platform-independent access to the patient record for authorized users within VUMC, at home, and at affiliated health care organizations.

BACKGROUND

The Medical ARchival System (MARS) was developed at the University of Pittsburgh Medical Center (UPMC) by John K. Vries and Russell J. Yount.[12] MARS stores massive amounts of data on clusters of UNIX workstations with high-performance magnetic disks. The system is designed to be scalable, and can keep a medical center's entire clinical and financial information on-line indefinitely. MARS provides full-text indexing of documents; any word in any document -- e.g., a discharge summary, a lab result such as a CBC, or an echocardiography report -- is indexed and searchable. This feature supports extremely powerful full-text queries, which can effectively search the entire patient record for any combination of words or phrases. The system uses parallel processing to provide excellent query and retrieval performance.

In March 1995, UPMC and VUMC signed a collaboration agreement which gave VUMC access to the MARS software. Implementation of MARS at Vanderbilt began at the same time. After an initial demonstration to hospital administrators and clinical service representatives in May '95, pilot testing began on five clinical units in October '95. MARS was released for housewide use at VUMC in January 1996. The system is now available on the 600 clinical workstations deployed through the Vanderbilt University Hospital and the Vanderbilt Clinic.

Implementing MARS consisted of several steps. First of all, we adapted the server side of the system, which runs on UNIX machines, to the VUMC environment.

0195—4210/96/$5.00 © 1996 AMIA, Inc.

This process took just a few days, and resulted in a fully functional, but unpopulated, data base. We then developed connections to the various VUMC data feeder systems, using a variety of transport mechanisms and protocols. At the same time, we began the development of a new user interface client.

The UPMC version of MARS includes a Microsoft Windows 3.1 interface, and the original character-only interface. The Windows interface is commonly used on the UPMC campus, and the character-only interface provides remote access via terminal emulators for modem- or Internet-based clients. The primary environmental difference between UPMC and Vanderbilt is that the VUMC clinical workstations run the IBM OS/2 operating system. In addition to this difference, the development of a new user interface client was dictated by our desire to give users immediate visual cues about which tests in VUMC's lab hierarchy, and which types of reports, are actually available for an individual patient. This information is presented by highlighting buttons. Clicking on a button gives users immediate access to a) the different types of report, and b) VUMC-specific spreadsheets of lab tests, from CBCs to microbiology culture reports. Fig. 1 shows a screen dump of the VUMC OS/2 MARS interface. A button with a gray background indicates that results of that type are available for that patient.

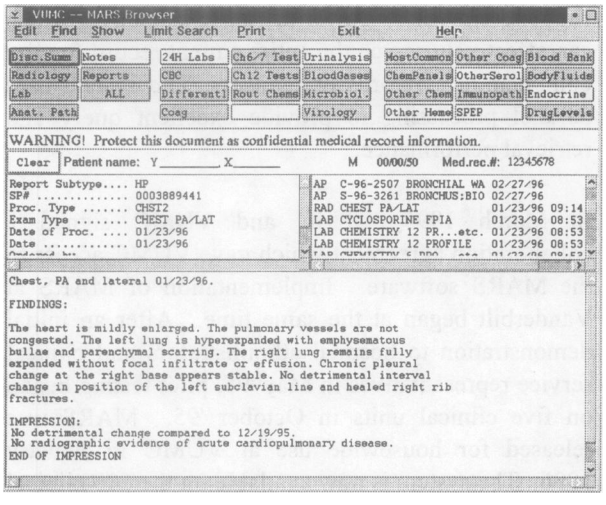

Figure 1 - VUMC OS/2 MARS client: sample screen. Identifying information has been changed to protect patient confidentiality.

Reactions to the MARS computerized patient record have been extremely positive. By July 1996, over 1,500 individual requests for MARS accounts had been approved (at that time, only physicians, case managers, and medical students on clinical rotations could access the system). All users must read and sign a detailed confidentiality agreement, which: describes the precautions necessary when handling confidential electronic records; explains that all user searches are logged permanently and audited for improper access; and specifies the penalties for misuse of electronic patient data -- up to, and including, dismissal from employment. Users conduct over 3,600 searches every day, and this number is growing rapidly.

REQUIREMENTS

While the availability of the OS/2 MARS client has satisfied the needs of clinical users in the hospital and clinics, it does not address areas without standardized VUMC workstations, such as faculty offices and research labs. The number of personal computers in such areas (estimated at 2,000) is greater than that of existing clinical workstations. This number includes both Intel-based and Macintosh models, running a variety of operating systems. Several physicians have also expressed interest in being able to access MARS records from home.

The variety of operating systems and hardware platforms in the VUMC community at large effectively precludes solutions which require client software to be developed for individual configurations. Therefore, the most practical approach is to develop World Wide Web[13] based applications, which can exploit the near-universal availability of Web browsers[14,15]. As long as security and confidentiality concerns are addressed appropriately, Web-based applications can also provide access outside the area covered by the VUMC backbone network.

Creation of a Web-based interface to MARS began in February 1996. In addition to platform independence, the new design included several requirements. First, it was necessary to retain the overall functionality of the MARS OS/2 client, such as the organization of the lab spreadsheets and the visual feedback about what types of reports are available for each patient. Second, we considered it important to be able to link the MARS client to additional documents such as on-line help, reference material, etc. Third, we wanted to experiment with a new, time-oriented

format showing users what reports are available. While the original OS/2 reports display (a scrollable window of all reports in reverse chronological order, shown on the right-hand side of Fig. 1) is adequate, we wanted to provide a more immediate view of how reports cluster together in time. The Web's suitability for rapid prototyping[14] facilitated experimentation with the new display.

Finally, users are increasingly requesting that non-textual information (such as EKG tracings and radiological images) be available as part of -- or in conjunction with -- the MARS patient record. The choice of the Web as the development platform clearly has the potential to address those multimedia requirements, as shown by projects such as the Image Engine.[15]

DESIGN OF THE WEB CLIENT

The main design challenge for the MARS Web client was to compensate for the connectionless nature of the HTTP protocol.[16] By definition, a connectionless protocol does not retain any state information between successive requests. However, the OS/2 client uses a substantial amount of state information to improve performance and give better feedback to the user. After a successful query, the client obtains a list of all the matching patient records, and determines which buttons should be highlighted to indicate what types of report are available. This information is cached in the client until the next query. Moreover, the OS/2 client caches all available lab results, and dynamically organizes them into a spreadsheet when the user clicks on a spreadsheet button. Finally, the OS/2 client keeps identifying information about the patient whose records are currently displayed. The amount of state information kept in the client averages about 150K bytes during a typical interactive session.

A naive implementation of a Web client would repeat the MARS search, obtain the same list of patient records, and retrieve documents every time the user clicked a button. In practice, this strategy would be hopelessly inefficient. The actual Web client splits the functionality into two components. The first component, requiring interim storage of data, is pushed back onto a server which sits between the client and the MARS database. This component creates a separate server process for each Web browser, and coordinates the interactions based on unique identifiers, similar to the approach described by Cimino et al[14]. To avoid unauthorized use, each process times out after a specified period of time, preventing access from HTML pages cached in a client workstation's Web browser. The second component is a set of Common Gateway Interface (CGI) scripts on the Web server that are invoked after a user action. The scripts receive data from the server process, format the data dynamically into new HTML documents, and return the HTML documents to the Web browser. Fig. 2 shows a diagram of the system, which is patterned after other systems that require a connection-oriented view (a common situation for Web-based interfaces to external databases).[17]

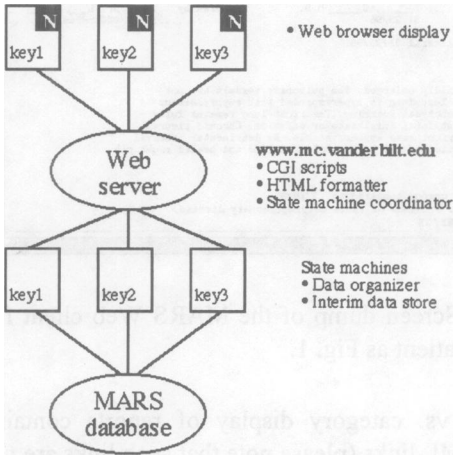

Figure 2: Architecture of the MARS Web client.

The scrollable list of reports in the OS/2 client has been replaced in the Web client by a 2-dimensional time vs. category display of all available reports, which was inspired by the time-oriented view in the THERESA™ system.[18] Each entry in the display indicates the total number of reports of a certain type for the given time period, as shown in Fig. 3; an asterisk indicates ten or more reports. Our display uses a variable time scale to show the entire known clinical history for a patient. Each column in the display corresponds to an appropriate unit of time, which can range from one day to one year. Like the THERESA view, our display of reports gives users an immediate understanding of when different types of clinical data were generated. This can be very effective to suggest periods of hospitalization (heavy generation of many report types, including discharge summaries) vs. outpatient episodes (isolated lab tests, echo reports, etc.).

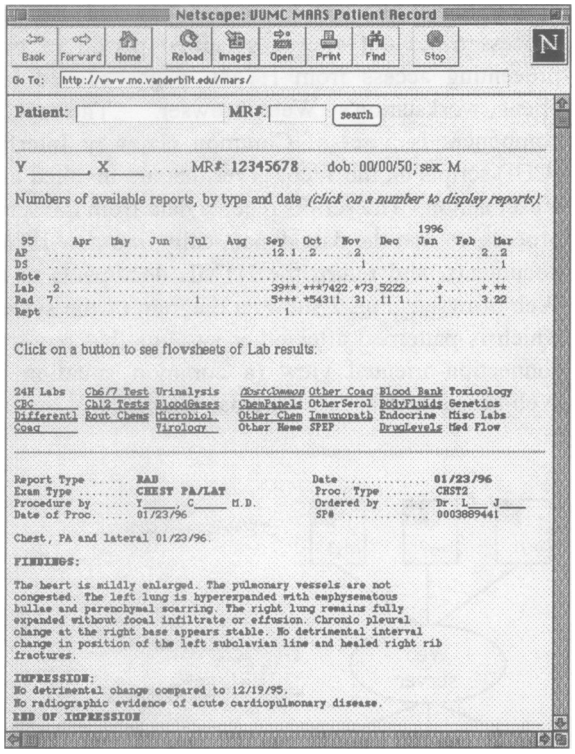

Figure 3: Screen dump of the MARS Web client for the same patient as Fig. 1.

The time vs. category display of reports contains active HTML links (please note that such links are not highlighted in Fig. 3, to prevent artifacts in the reproduction of the sample screen). Clicking on a link (i.e., a number) in a row retrieves and displays the reports of the given type for that period of time. This allows the user, for example, to click on the number "1" in the Radiology row in Fig. 3 to display the report of a PA/Lateral chest X-ray that was taken on January 23, 1996. When the number is greater than one, all reports for that time period are formatted and displayed in one HTML document.

Once an HTML document is displayed by a Web browser, it is impossible to make incremental changes to the display (unless not-yet standardized HTML features such as Netscape Navigator frames [19] or the JAVA language are used). Therefore, the server generates a new HTML document every time the display needs to change -- for example, because the user clicks on a button to display a spreadsheet. Despite this problem, the performance of the Web client is good and quite similar to that of the OS/2 client. Because of the need to redisplay the entire document every time, however, the Web interface tends to cause more image flicker than the OS/2 interface.

As others have noted,[14] implementing applications that rely on HTML documents is not time-consuming. The implementation of the Web-based MARS client had taken about 2 weeks FTE as of the middle of March 1996, producing the same functionality as that of the original OS/2 interface. Of course, most of the logic built for the original interface was simply reused, by relying on the same compiled library for both interfaces. Much of the new development time was in deciding what HTML constructs would work best, and in implementing the set of CGI scripts.

DISCUSSION

Our current plans are to keep the OS/2 client as the main interface to MARS from the clinical workstations, because clinical users are already well trained in its use and its look-and-feel. Users on all other workstations will be directed to the Web client. In particular, users who have OS/2-based non-standard workstations will also use the Web client. This will eliminate the need to support the OS/2 client on machines that are not as tightly controlled as the VUMC clinical workstations.

Final plans for allowing outside access to the MARS records via the Web have not yet been finalized as of March 1996. Currently, only requests that originate from the VUMC network domain (and which undergo the normal authentication via user ID and password) are allowed. We are in the process of evaluating the interaction between the firewall and outside Web access, and also investigating the use of secure HTTP technology as a way of preventing unauthorized access to confidential data while in transit between VUMC and remote sites.

In spite of the technical challenges inherent in World Wide Web technology, its advantages as a platform-independent, multimedia transport vehicle are substantial. We see its universal availability and ease of development as a key factor in supporting VUMC's strategic alliances with other health care organizations, and as a way of addressing the growing need to exchange clinically relevant data among institutions that have traditionally been isolated.

Acknowledgments

The research described in this document was supported in part by the National Library of Medicine IAIMS grant no. 5 G08 LM05443-04. We thank Tommy Williams for many constructive suggestions about programming with CGI scripts. We also thank Randolph Miller, Suzanne Kafantaris, and Steve Brown for helpful comments on the manuscript.

References

1. Institute of Medicine. *The Computer-based Patient Record: An Essential Technology for Health Care.* Washington, DC: National Academy Press, 1991.

2. Johnson SB, Cimino JJ, Friedman C, Hripcsak G, Clayton PD. Using metadata to integrate medical knowledge in a clinical information system. In: *Proc 14th Ann Symp Comp Applic Med Care.* IEEE Computer Society Press, 1990:340-344.

3. Board of Directors of the American Medical Informatics Association. Standards for Medical Identifiers, Codes, and Messages Needed to Create an Efficient Computer-Stored Medical Record. *J Am Med Informatics Assoc.* 1994; 1: 1-7.

4. Greenes RA, Papillardo AN, Marble CW, Barnett GO. Design and implementation of a clinical data management system. *Comput Biomed Res.* 1969; 2:469-485.

5. Stead WW, Hammond WE. Computer-based medical records: the centerpiece of TMR. *MD Comput.* 1988;5(5):48-62.

6. Barnett GO. Massachusetts General Hospital computer system. In: Collen MF (ed). *Hospital Computer Systems.* New York: Wiley & Sons, 1974:517-545.

7. Warner HR, Olmsted CM, Rutherford BD. HELP -- a program for medical decision making. *Comput Biomed Res.* 1972;5:65-74.

8. McDonald CJ. Protocol-based computer reminders, the quality of care and the nonperfectibility of man. *N Engl J Med.* 1976;295:1351-1355.

9. Pryor TA, Gardner RM, Clayton PD, Warner HR. The HELP system. *J Med Sys.* 1983;7:87-101.

10. Perry M, Myers CE. Computer cost prompting as a determinant of hospital drug prescribing. *Proc 18th Annu Am Soc Hosp Pharm (AHSP) Mid-year Clinical Meeting.* Atlanta, GA, 1983.

11. Tierney WM, McDonald CJ, Martin DK, Rogers MP. Computerized display of past test results: effect on outpatient testing. *Ann Intern Med.* 1987;107:569-574.

12. Yount RJ, Vries JK, Councill CD. The Medical ARchival System: an Information Retrieval System Based on Distributed Parallel Processing. *Inf Processing & Management.* 1991;27(4):1-11.

13. Berners-Lee T, Cailliau R, Luotonen A, Nielsen HF, Secret A. The World Wide Web. *Comm of the ACM.* 1994; 37(8):76-82.

14. Cimino JJ, Socratous SA, Grewal R. The Informatics Superhighway: Prototyping on the World Wide Web. In: *Proc 19th Annu Symp Comp Appl Med Care.* 1995:111-115.

15. Lowe HJ, Walker WK, Vries JK. Using Agent-Based Technology to Create a Cost Effective, Integrated, Multimedia View of the Electronic Medical Record. In: Gardner RM (ed). *Proc 19th Annu Symp Comp Appl Med Care.* 1995:441-444.

16. Internet Engineering Task Force. *HyperText Markup Language Specification Version 3.0.* IETF, 1995.

17. Willard KE, Hallgren JH, Sielaff B, Connelly DP. The Deployment of a World Wide Web (W3) Based Medical Information System. In: Gardner RM (ed). *Proc 19th Annu Symp Comp Appl Med Care.* 1995:771-775.

18. Camp HN. Technical challenges, past and future, in implementing THERESA - a one million patient, one billion item computer-based patient record and decision support system. *Proc Health Care Information Infrastructure.* SPIE, Bellingham, WA, 1995.

19.http://developer.netscape.com/platform/html_compilation/forms.htm